

Diversity of preferences can increase collective welfare in sequential exploration problems

PANTELIS P. ANALYTIS, Cornell University

HRVOJE STOJIC, University College London

ALEXANDROS GELASTOPOULOS, Boston University

MEHDI MOUSSAÏD, Max Planck Institute for Human Development

1. INTRODUCTION

In search engines, online marketplaces and other human–computer interfaces large collectives of individuals sequentially interact with numerous alternatives of varying quality. In these contexts, individual trial and error (exploration) is crucial for uncovering novel high-quality items or solutions, but entails a high cost for individual agents [Frazier et al. 2014]. Self-interested decision makers, we will show, are often better off imitating the choices of individuals who have already incurred the costs of exploration. Although imitation makes sense at the individual level, it deprives the group of additional information that could have been gleaned by individual explorers [Rogers 1988]. Under these grim circumstances, certain non-monetary mechanisms can keep imitation forces in check and allow the collective to reap some of the benefits of the independent collection of information. For example, in simultaneous exploration problems, a natural equilibrium evolves between explorers and imitators [Conlisk 1980; Kameda and Nakanishi 2002]. Further, in some collective exploration settings, barriers to communication such as a sparser communication network among individuals can prove beneficial at the collective level. They encourage people to explore more, thus supplying useful information to the group [Fang et al. 2010; Lazer and Friedman 2007; Mason et al. 2008; Toyokawa et al. 2014].

Diversity is known to be a blessing for groups and collectives, as they can leverage the wealth of information possessed by different individuals [Conradt et al. 2013; Davis-Stober et al. 2014; Müller-Trede et al. 2017] or take advantage of the complementarities between group members to solve complex problems [Clearwater et al. 1991; Hong and Page 2004]. Could some preference diversity be beneficial in problems where collectives sequentially explore numerous alternatives, and thus despite reducing the immediate value of social learning lead to an increase in collective welfare?

2. THE MODEL

We use a search model to study the problem, as it exemplifies the trade-off between exploration and exploitation [March 1991]. We assume that there are N alternatives (or products) X_1, \dots, X_N in a market populated by M agents A_1, \dots, A_M . The alternatives have an objective utility component u_o , which is identical for all agents, and a subjective component u_s , which is agent specific. The objective component u_o of each alternative is a draw from an independent and identically distributed (iid) random variable, normally distributed with mean μ_o and variance σ_o^2 . The subjective component u_s of each alternative is iid normal with mean μ_s and variance σ_s^2 . The overall utility of an alternative n is then a sum of two draws from iid normal variables, $u_n = u_{no} + u_{ns}$, which itself is an iid normal variable with variance $\sigma^2 = \sigma_o^2 + \sigma_s^2$.

The agents encounter alternatives sequentially and can learn the utility of an alternative u_n only by sampling it and paying a search cost c . The agents can sample as many alternatives as desired, but they can choose only one of the alternatives they have sampled. An example of sampling behavior would be

skimming through book recommendations at Amazon or examining the results provided by a search engine. We assume that the search cost is constant and that there is no post-sampling uncertainty—by sampling the alternative, the agents learn its true utility. An agent’s return depends on the utility of the best alternative discovered so far and the search cost, $\max(u_1, u_2, \dots, u_n) - n \times c$. The returns from sampling one more alternative can be formulated as:

$$E(u_n - \max(u_1, u_2, \dots, u_{n-1}) | u_n > \max(u_1, u_2, \dots, u_{n-1})) - c$$

When viewed on the level of a single agent, this is a classical optimal stopping problem—as studied extensively in statistics and economics [DeGroot 1970; Analytis et al. 2015]—in which after examining each new alternative the agent decides whether to sample further or to stop search. For random search, the problem has an optimal solution, which can be expressed in the form of a stopping threshold.

In our model, agents do not search randomly. Instead, they follow public popularity information generated by choices made by previous searchers in the market. After the first agent, a second or m^{th} agent is asked to make its choice. Now the agent can observe the sum of choices made by previous searchers in the form of a popularity vector $P = \{P_1, \dots, P_N\}$ that records the choices of alternatives X_1, \dots, X_N and is updated whenever an individual makes a choice. Examples of such popularity information would be product sales, number of song downloads, or number of article citations. The agents sample the alternatives X_1, X_2, \dots, X_N in decreasing order of popularity $P_1 \geq P_2 \geq \dots \geq P_N$, where popularity is defined as the number of times that an alternative has been selected in the past. When there is a tie, the agents choose which one to sample next at random. We call the order in which the alternatives are searched the *search path*. The agents stop search when they encounter an alternative u_n with utility higher than a threshold T ; otherwise, they continue sampling until the alternatives are exhausted. The proposed model captures the behavioral regularities of search in online interfaces [Joachims et al. 2005; Craswell et al. 2008] and bolsters them with a fully fleshed-out utility framework.

3. SIMULATIONS AND RESULTS

We simulated markets consisting of 1000 agents that decided sequentially which of 100 alternatives to choose, after searching the market according to public popularity information. We systematically varied the diversity of preferences in the market, d , by setting the subjective variance of each individual agent σ_s^2 equal to $d \in \{0, 0.1, \dots, 0.9, 1\}$ and the variance of the objective utility σ_o^2 equal to $1 - d$. We drew subjective utilities for each agent independently, whereas the draw of the objective utilities was the same for all agents. Next, we varied the cost of search, c , that agents have to pay when sampling an alternative, $c = \{1/2^2, 1/2^3, 1/2^4, 1/2^5, 1/2^6, 1/2^7, 1/2^8\}$. These costs correspond to the threshold parameters $T = \{0.34, 0.78, 1.15, 1.47, 1.76, 2.03, 2.28\}$; these parameters are optimal for a random search model [Chow et al. 1971], which our agents used despite searching according to the popularity ranking. This yields $11 \times 7 = 77$ markets, each of which was simulated 1000 times to obtain stable averages.¹

In Figure 1 (left), we illustrate the average utility and total cost for the agents in the market, focusing on the second higher search cost case $c = 1/2^3$. In the maximum diversity case, $d = 1$, the environment is recreated randomly for each agent. In this case, the popularity information is completely uninformative and welfare outcomes are equal to those in the market where agents search randomly. At $d = 1$, agents incur the highest average total costs and earn the lowest average net utility from the market, even though the thresholds they employ are optimally set. At $d = 0$, the alternative chosen by the first agent satisfies all the subsequent agents in the population, leading to herding. In these markets, the cost of search is minimal, and the obtained net utility is much higher than in the random search market. Because the first agent may settle on a very good or a simply passable alternative, $d = 0$ implies

¹The code for the simulations can be openly accessed at <https://osf.io/up5pe/>

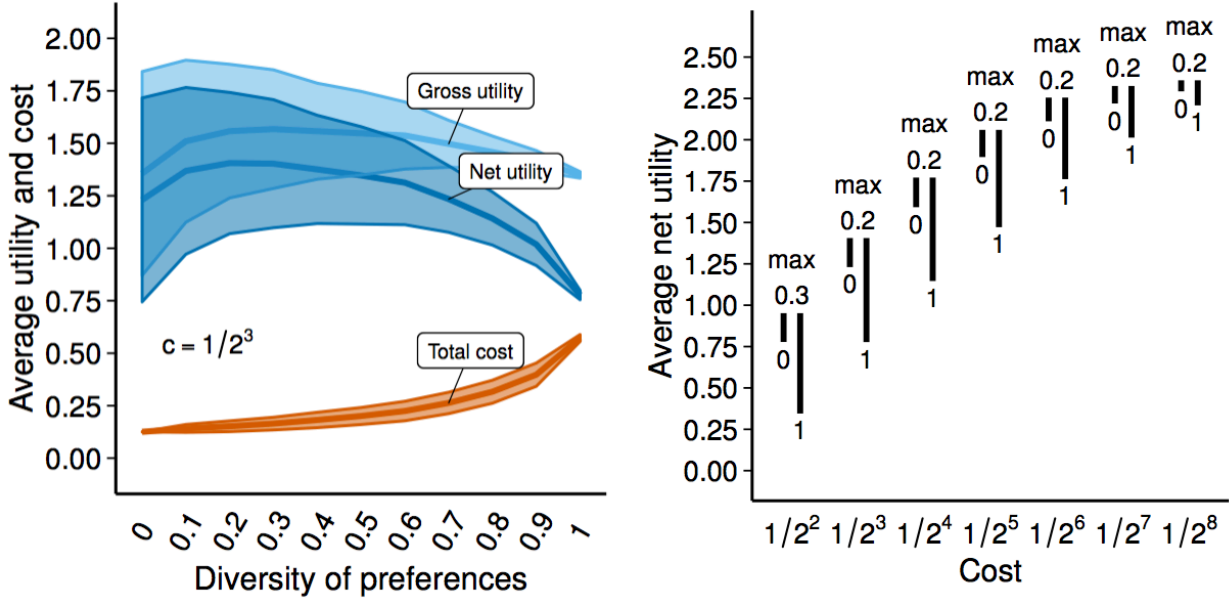


Fig. 1: The average net utility in the market is highest at intermediate levels of preference diversity. The lines represent the average gross and net utility of the selected alternative and the average total cost of search at a search cost of $1/2^3$. The ribbons represent the variability in 1000 repetitions of the simulation. The beneficial effect of diversity increases with the search cost. For each cost level, we illustrate the difference between the level of diversity with the maximum net utility and conditions with no diversity, $d = 0$, and maximum diversity, $d = 1$.

more variability in the average utility in the market. This is illustrated by the ribbons in Figure 1 (left). Although the search costs are minimized, the economy does not maximize the collective welfare of the agents involved. There is a robust non-monotonic relationship between d and the average net utility—the highest level of utility is consistently achieved for $d \approx 0.2$, whereas at high costs even markets with $d = 0.5$ on average do better than markets with no diversity whatsoever. Although the average agent incurs a higher average cost of search in markets with $d > 0$ than in a market with $d = 0$, this loss is outweighed by the benefits of choosing alternatives with higher gross utility. Figure 1 (right) illustrates that the non-monotonic relation between diversity and collective utility persists even at much lower costs of search, although the average differences between conditions decrease.

What is the mechanism that leads to an increase in average net utility in markets with some diversity of preferences? In a market with no diversity whatsoever, everybody herds on the option selected by the first individual. Occasionally, this turns out to be an excellent alternative. More often than not, however, it is located just off the satisficing threshold. The collective welfare is initial outcome-dependent [Page 2006]. In contrast, in markets with some diversity of preferences, agents do not always appreciate the choices of previous individuals and they venture further until they encounter a satisficing option. With each choice, the *search path* followed by the agents improves. Eventually, alternatives with higher objective utility components are more likely to occupy the first positions in the *search path*. Thus, agents deciding later on in time possess a useful popularity signal, require relatively few steps to find good alternatives, and tend to settle on better outcomes (Figure 1, gross utility). In essence, some diversity of preferences allows the agents to benefit from the thick informational externalities that are inherent in sequential social interactions. By searching and deciding, the agents unintentionally provide useful information to the individuals deciding after them.

REFERENCES

- Pantelis P. Analytis, Hrvoje Stojic, and Mehdi Moussaïd. 2015. The collective dynamics of sequential search in markets for cultural products. *Santa Fe Institute Working Paper No. 2015-06-023* (2015).
- Yuan Shih Chow, Herbert Robbins, and David Siegmund. 1971. *Great expectations: The theory of optimal stopping*. Boston, MA: Houghton Mifflin Boston.
- Scott H. Clearwater, Bernardo A. Huberman, and Tad Hogg. 1991. Cooperative solution of constraint satisfaction problems. *Science* 254 (1991), 1181–1183.
- John Conlisk. 1980. Costly optimizers versus cheap imitators. *Journal of Economic Behavior & Organization* 1, 3 (1980), 275–293.
- Larissa Conradt, Christian List, and Timothy J. Roper. 2013. Swarm intelligence: When uncertainty meets conflict. *The American Naturalist* 182, 5 (2013), 592–610.
- Nick Craswell, Onno Zoeter, Michael Taylor, and Bill Ramsey. 2008. An experimental comparison of click position-bias models. In *Proceedings of the 2008 International Conference on Web Search and Data Mining*. ACM, Palo Alto, CY, 87–94.
- Clinton P. Davis-Stober, David V. Budescu, Jason Dana, and Stephen B. Broomell. 2014. When is a crowd wise? *Decision* 1, 2 (2014), 97–101.
- Morris H. DeGroot. 1970. *Optimal statistical decisions*. New York, NU: McGraw-Hill.
- Christina Fang, Jeho Lee, and Melissa A. Schilling. 2010. Balancing exploration and exploitation through structural design: The isolation of subgroups and organizational learning. *Organization Science* 21, 3 (2010), 625–642.
- Peter Frazier, David Kempe, Jon Kleinberg, and Robert Kleinberg. 2014. Incentivizing exploration. In *Proceedings of the fifteenth ACM conference on Economics and computation*. ACM, 5–22.
- Lu Hong and Scott E. Page. 2004. Groups of diverse problem solvers can outperform groups of high-ability problem solvers. *Proceedings of the National Academy of Sciences of the United States of America* 101, 46 (2004), 16385–16389.
- Thorsten Joachims, Laura Granka, Bing Pan, Helene Hembrooke, and Geri Gay. 2005. Accurately interpreting clickthrough data as implicit feedback. In *Proceedings of the 28th Annual International ACM SIGIR Conference on Research and Development in Information Retrieval*. ACM, Salvador, Brazil, 154–161.
- Tatsuya Kameda and Daisuke Nakanishi. 2002. Cost–benefit analysis of social/cultural learning in a nonstationary uncertain environment: An evolutionary simulation and an experiment with human subjects. *Evolution and Human Behavior* 23, 5 (2002), 373–393.
- David Lazer and Allan Friedman. 2007. The network structure of exploration and exploitation. *Administrative Science Quarterly* 52, 4 (2007), 667–694.
- James G March. 1991. Exploration and exploitation in organizational learning. *Organization Science* 2, 1 (1991), 71–87.
- Winter A. Mason, Andy Jones, and Robert L. Goldstone. 2008. Propagation of innovations in networked groups. *Journal of Experimental Psychology: General* 137, 3 (2008), 422–433.
- Johannes Müller-Trede, Shoham Choshen-Hillel, Meir Barneron, and Ilan Yaniv. 2017. The wisdom of crowds in matters of taste. *Management Science* (2017). DOI: <http://dx.doi.org/10.1287/mnsc.2016.2660>
- Scott E Page. 2006. Path Dependence. *Quarterly Journal of Political Science* 1 (2006), 87–115.
- Alan R. Rogers. 1988. Does biology constrain culture? *American Anthropologist* 90, 4 (1988), 819–831.
- Wataru Toyokawa, Hye-rin Kim, and Tatsuya Kameda. 2014. Human collective intelligence under dual exploration–exploitation dilemmas. *PLOS ONE* 9, 4 (2014), e95789.